arXiv:1512.01829v1 [cs.DS] 6 Dec 2015

On Routing Disjoint Paths in Bounded Treewidth Graphs

Alina Ene* Matthias Mnich† Marcin Pilipczuk‡ Andrej Risteski§

Abstract

We study the problem of routing on disjoint paths in bounded treewidth graphs with both edge and node capacities. The input consists of a capacitated graph G and a collection of k source-destination pairs $\mathcal{M} = \{(s_1, t_1), \ldots, (s_k, t_k)\}$. The goal is to maximize the number of pairs that can be routed subject to the capacities in the graph. A routing of a subset $\mathcal{M}'$ of the pairs is a collection $\mathcal{P}$ of paths such that, for each pair $(s_i, t_i) \in \mathcal{M}'$, there is a path in $\mathcal{P}$ connecting s_i to t_i. In the Maximum Edge Disjoint Paths (MaxEDP) problem, the graph G has capacities $\mathrm{cap}(e)$ on the edges and a routing $\mathcal{P}$ is *feasible* if each edge e is in at most $\mathrm{cap}(e)$ of the paths of $\mathcal{P}$. The Maximum Node Disjoint Paths (MaxNDP) problem is the node-capacitated counterpart of MaxEDP.

In this paper we obtain an $\mathcal{O}(r^3)$ approximation for MaxEDP on graphs of treewidth at most r and a matching approximation for MaxNDP on graphs of pathwidth at most r. Our results build on and significantly improve the work by Chekuri et al. [ICALP 2013] who obtained an $\mathcal{O}(r \cdot 3^r)$ approximation for MaxEDP.

1 Introduction

In this paper, we study disjoint paths routing problems on bounded treewidth graphs. In this setting, we are given an undirected capacitated graph G and a collection of source-destination pairs $\mathcal{M} = \{(s_1, t_1), (s_2, t_2), \ldots, (s_k, t_k)\}$. The goal is to select a maximum-sized subset $\mathcal{M}' \subseteq \mathcal{M}$ of the pairs that can be *routed* subject to the capacities in the graph. More precisely, a routing of $\mathcal{M}'$ is a collection $\mathcal{P}$ of paths such that, for each pair $(s_i, t_i) \in \mathcal{M}'$, there is a path in $\mathcal{P}$ connecting s_i to t_i. In the Maximum Edge Disjoint Paths (MaxEDP) problem, the graph G has capacities $\mathrm{cap}(e)$ on the edges and a routing $\mathcal{P}$ is *feasible* if each edge e is in at most $\mathrm{cap}(e)$ of the paths of $\mathcal{P}$. The Maximum Node Disjoint Paths (MaxNDP) problem is the node-capacitated counterpart of MaxEDP.

Disjoint paths problems are fundamental problems with a long history and significant connections to optimization and structural graph theory. The decision versions of MaxEDP and MaxNDP ask whether all of the pairs can be routed subject to the capacities. Karp showed that, when the number of pairs is part of the input, the decision problems are NP-complete (the node disjoint paths is part of Karp's original list of NP-complete problems [18]). In undirected graphs, MaxEDP and MaxNDP are solvable in polynomial time when the number of pairs is a fixed constant; this is a very deep result of Robertson and Seymour [24] that builds on several fundamental results in structural graph theory from their graph minors project.

In this paper, we consider the optimization problems MaxEDP and MaxNDP when the number of pairs are part of the input. These problems are NP-hard and the main focus in this paper is on approximation algorithms for these problems in bounded treewidth graphs. Although they may appear to be quite specialized at first, MaxEDP and MaxNDP on capacitated graphs of small treewidth capture a surprisingly rich class of problems; in fact, as shown by Garg, Vazirani, and Yannakakis [16], these problems are quite interesting and general even on trees.

*Department of Computer Science and DIMAP, University of Warwick. A.Ene@dcs.warwick.ac.uk

†Institut für Informatik, Universität Bonn. mmnich@uni-bonn.de. Supported by ERC Starting Grant 306465 (BeyondWorstCase).

‡Institute of Informatics, University of Warsaw. malcin@mimuw.edu.pl. Research done while at University of Warwick, partially supported by DIMAP and by Warwick-QMUL Alliance in Advances in Discrete Mathematics and its Applications.

§Department of Computer Science, Princeton University. risteski@princeton.edu

MaxEDP and MaxNDP have received considerable attention, leading to several breakthroughs both in terms of approximation algorithms and hardness results. MaxEDP is APX-hard even in edge-capacitated trees [16], whereas the decision problem is trivial on trees; thus some of the hardness of the problem stems from having to select a subset of the pairs to route. Moreover, by subdividing the edges, one can easily show that MaxNDP generalizes MaxEDP in capacitated graphs. However, node capacities pose several additional technical challenges and extending the results for MaxEDP to the MaxNDP setting is far from immediate even in restricted graph classes and our understanding of MaxNDP is more limited.

In general graphs, the best approximation for MaxEDP and MaxNDP is an $\mathcal{O}(\sqrt{n})$ approximation [6, 19], where n is the number of nodes, whereas the best hardness for undirected graphs is only $\Omega((\log n)^{1/2-\epsilon})$ [2]. Bridging this gap is a fundamental open problem that seems quite challenging at the moment. There have been several breakthrough results on a relaxed version of these problems where congestion is allowed[1]. This line of work has culminated with a polylog(n) approximation with congestion 2 for MaxEDP [13] and congestion 51 for MaxNDP [5]. In addition to the routing results, this work has led to several significant insights into the structure of graphs with large treewidth and to several surprising applications [4].

Most of the results for routing on disjoint paths use a natural multi-commodity flow relaxation as a starting point. A well-known integrality gap instance due to Garg et al. [16] shows that this relaxation has an integrality gap of $\Omega(\sqrt{n})$, and this is the main obstacle for improving the $\mathcal{O}(\sqrt{n})$ approximation in general graphs. The integrality gap example is an instance on an $n \times n$ grid that exploits a topological obstruction in the plane that prevents a large integral routing (see Fig. 2). Since an $n \times n$ grid has treewidth $\Theta(\sqrt{n})$, it suggests the following natural and tantalizing conjecture that was asked by Chekuri et al. [9].

Conjecture 1 ([9])**.** *The integrality gap of the standard multi-commodity flow relaxation for* MaxEDP *(and* MaxNDP*) is* $\Theta(r)$ *with congestion* 1*, where* r *is the treewidth of the graph.*

Recently, Chekuri, Naves, and Shepherd [10] showed that MaxEDP admits an $\mathcal{O}(r \cdot 3^r)$ approximation on graphs of treewidth at most r. This is the first approximation for the problem that is independent of n and k, and the first step towards resolving the conjecture. One of the main questions left open by the work of Chekuri et al. [10] — that was explicitly asked by them — is whether the exponential dependency on the treewidth is necessary. In this paper, we address this question and we make a significant progress towards resolving Conjecture 1.

Theorem 1.1. *The integrality gap of the multi-commodity flow relaxation is* $\mathcal{O}(r^3)$ *for* MaxEDP *in edge-capacitated undirected graphs of treewidth at most* r*. Moreover, there is a polynomial time algorithm that, given a tree decomposition of* G *of width at most* r *and a fractional solution to the relaxation of value* OPT*, it constructs an integral routing of size* $\Omega(\mathrm{OPT}/r^3)$*.*

As mentioned above, MaxNDP in node-capacitated graphs is more general than MaxEDP and it poses several additional technical challenges. In this paper, we give an $\mathcal{O}(r^3)$ approximation for MaxNDP on graphs of pathwidth at most r with arbitrary node capacities. This is the first result for MaxNDP that is independent of n and it improves the $\mathcal{O}(r \log r \log n)$ approximation of Chekuri et al. [9].

Theorem 1.2. *The integrality gap of the multi-commodity flow relaxation is* $\mathcal{O}(r^3)$ *for* MaxNDP *in node-capacitated undirected graphs of pathwidth at most* r*. Moreover, there is a polynomial time algorithm that, given a path decomposition of* G *of width at most* r *and a fractional solution to the relaxation of value* OPT*, it constructs an integral routing of size* $\Omega(\mathrm{OPT}/r^3)$*.*

The study of routing problems in bounded treewidth graphs is motivated not only by the goal of understanding the integrality gap of the multi-commodity flow relaxation but also by the broader goal of giving a more refined understanding of the approximability of routing problems in undirected graphs. Andrews et al. [2] have shown that MaxEDP and MaxNDP in general graphs cannot be approximated within a factor better than $(\log n)^{\Omega(1/c)}$ even if we allow a constant congestion $c \geq 1$. Thus in order to obtain constant factor

[1] A collection of paths has an *edge* (resp. *node*) *congestion* of c if each edge (resp. node) is in at most $c \cdot \mathrm{cap}(e)$ (resp. $c \cdot \mathrm{cap}(v)$) paths.

approximations one needs to use additional structure. However, this seems challenging with our current techniques and there are only a handful of results in this direction.

One of the main obstacles for obtaining constant factor approximations for disjoint paths problems is that most approaches rely on a powerful pre-processing step that reduces an arbitrary instance of MaxEDP/MaxNDP to a much more structured instance in which the terminals[2] are *well-linked*. This reduction is achieved using the well-linked decomposition technique of Chekuri, Khanna, and Shepherd [7], which necessarily leads to an $\Omega(\log n)$ loss even in very special classes of graphs such as bounded treewidth graphs. Chekuri, Khanna, and Shepherd [8] showed that the well-linked decomposition framework can be bypassed in planar graphs, leading to a $\mathcal{O}(1)$ approximation for MaxEDP with congestion 4 (the congestion was later improved by Séguin-Charbonneau and Shepherd [26] from 4 to 2). This result suggests that it may be possible to obtain constant factor approximations with constant congestion for much more general classes of graphs. In particular, Chekuri et al. [10] conjecture that this is the case for the class of all minor-free graphs.

Conjecture 2 ([10])**.** *Let $\mathcal{G}$ be any proper minor-closed family of graphs. Then the integrality gap of the multi-commodity flow relaxation for* MaxEDP *is at most a constant $c_{\mathcal{G}}$ when congestion* 2 *is allowed.*

A natural approach is to attack Conjecture 2 using the structure theorem for minor-free graphs given by Robertson and Seymour [23, 25] that asserts that every such graph admits a tree decomposition where the size of every adhesion (the intersection of neighboring bags) is bounded, and after turning the adhesions into cliques, every bag induces a structurally simpler graph: one of bounded genus, with potentially a bounded number of apices and vortices. Thus in some sense, in order to resolve Conjecture 2, one needs to understand the base graph class (bounded genus graphs with apices and vortices) and how to tackle bounded width tree decompositions.

The recent work of Chekuri et al. [10] has made a significant progress toward resolving Conjecture 2 by providing a toolbox for the latter issue, and the only ingredient that is still missing is an algorithm for planar and bounded genus graphs with a constant number of vortices (in the disjoint paths setting, apices are very easy to handle). However, one of the main drawbacks of their approach is that it leads to approximation guarantees that are *exponential* in the treewidth. Our work strengthens the approach of Chekuri et al. [10] and it gives a much more graceful *polynomial* dependence in the approximation ratio.

Theorem 1.3. *Let $\mathcal{G}$ be a minor-closed class of graphs such that the integrality gap of the multi-commodity flow relaxation is α with congestion β. Let $\mathcal{G}_\ell$ be the class of graphs that admit a tree decomposition where, after turning all adhesions into cliques, each bag induces a graph from $\mathcal{G}$, and each adhesion has size at most ℓ. Then the integrality gap of the relaxation for the class $\mathcal{G}_\ell$ is $\mathcal{O}(\ell^3) \cdot \alpha$ with congestion $\beta + 3$.*

We also revisit the well-linked decomposition framework of Chekuri et al. [7] and we ask whether the $\Omega(\log n)$ loss is necessary for very structured graph classes. For bounded treewidth graphs, we give a well-linked decomposition framework that reduces an arbitrary instance of MaxEDP to node-disjoint instances of MaxEDP that are *well-linked*. The loss in the approximation for our decomposition is only $\mathcal{O}(r^3)$, which improves the guarantee of $\mathcal{O}(\log r \log n)$ from Chekuri et al. [9] when r is much smaller than n.

It is straightforward to obtain the improved well-linked decomposition from our algorithm for MaxEDP. Nevertheless, we believe it is beneficial to have such a well-linked decomposition, given that well-linked decompositions are one of the technical tools at the heart of the recent algorithms for routing on disjoint paths, integral concurrent flows [3], and flow and cut sparsifiers [12]. In particular, we hope that such a well-linked decomposition will have applications to finding flow and cut sparsifiers with Steiner nodes for bounded treewidth graphs. A sparsifier for a graph G with k source-sink pairs is a significantly smaller graph H containing the terminals (and potentially other vertices, called Steiner nodes) that approximately preserves multi-commodity flows or cuts between the terminals. Such sparsifiers have been extensively studied and several results are known both in general graphs and in bounded treewidth graphs (see Andoni et al. [1] and references therein).

A different question one could ask for problems in bounded treewidth graphs is whether additional computational power beyond polynomial-time running time can help with MaxEDP or MaxNDP. It is a standard exercise to design an $n^{\mathcal{O}(r)}$-time dynamic programming algorithm (i.e., polynomial for every constant r) for

[2]The vertices participating in the pairs $\mathcal{M}$ are called *terminals*.

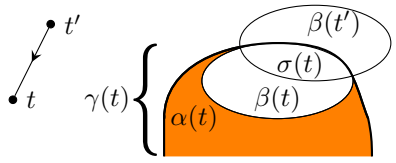

Figure 1: Notations used for a node t with parent t' in a tree decomposition $(\mathcal{T}, \beta)$.

MaxNDP in uncapacitated graphs of treewidth r, while the aforementioned results on hardness of MaxEDP in capacitated trees [16] rule out similar results for capacitated variants. Between the world of having r as part of the input, and having r as a fixed constant, lies the world of *parameterized complexity*, that asks for algorithms (called *fixed-parameter algorithms*) with running time $f(r) \cdot n^c$, where f is any computable function, and c is a constant independent of the parameter. It is natural to ask whether allowing such running time can lead to better approximation algorithms. As a first step towards resolving this question, we show a hardness for MaxNDP parameterized by *treedepth*, a much more restrictive graph parameter than treewidth (cf. [22]).

Theorem 1.4. MaxNDP *parameterized by the treedepth of the input graph is* $W[1]$*-hard, even with unit capacities.*

Consequently, the existence of an *exact* fixed-parameter algorithm is highly unlikely. We remark that our motivation for the choice of treedepth as a parameter stems also from the observation that a number of algorithms using the Sherali-Adams hierarchy to approximate a somewhat related problem of Nonuniform Sparsest Cut in bounded treewidth graphs [11, 17] in fact implicitly uses a rounding scheme based on treedepth rather than treewidth.

Paper organization. The rest of this paper is organized as follows. In Sect. 2, we formally define the problems and the multi-commodity flow relaxation. In Sect. 3, we give our algorithm for MaxEDP in bounded treewidth graphs and prove Theorems 1.1 and 1.3. In Sect. 4, we extend our MaxEDP algorithm to MaxNDP in bounded pathwidth graphs, and prove Theorem 1.2. We give a well-linked decomposition for edge-capacitated graphs in Sect. 5. Finally, we prove Theorem 1.4 in Sect. 6.

2 Preliminaries

Tree and path decompositions. In this paper all tree decompositions are rooted; that is, a tree decomposition of a graph G is a pair $(\mathcal{T}, \beta)$ where $\mathcal{T}$ is a rooted tree and $\beta : V(\mathcal{T}) \to 2^{V(G)}$ is a mapping satisfying the following properties: (i) for every $e \in E(G)$, there exists a node $t \in V(\mathcal{T})$ with $e \subseteq \beta(t)$; (ii) for every $v \in V(G)$ the set $\{t : v \in \beta(t)\}$ is nonempty and connected in $\mathcal{T}$.

For a node $t \in V(\mathcal{T})$, we call the set $\beta(t)$ the *bag* at node t, while for an edge $st \in E(\mathcal{T})$, the set $\beta(t) \cap \beta(s)$ is called an *adhesion*. For a non-root node $t \in V(\mathcal{T})$, by $\mathrm{parent}(t)$ we denote the parent of t, and by $\sigma(t) := \beta(t) \cap \beta(\mathrm{parent}(t))$ the adhesion on the edge to the parent of t, called henceforth *the parent adhesion*; for the root node $t_0 \in V(\mathcal{T})$ we put $\sigma(t_0) = \emptyset$. For two nodes $s, t \in V(\mathcal{T})$, we denote by $s \preceq t$ if s is a descendant of t, and put $\gamma(t) := \bigcup_{s \preceq t} \beta(s)$, $\alpha(t) := \gamma(t) \setminus \sigma(t)$, and $G(t) := G[\gamma(t)] \setminus E(G[\sigma(t)])$.

A *torso* at node t is a graph obtained from $G[\beta(t)]$ by turning every adhesion for an edge incident to t into a clique.

We say that (A, B) is a *separation* in G if $A \cup B = V(G)$ and there does not exist an edge of G with an endpoint in $A \setminus B$ and the other endpoint in $B \setminus A$. We use the following well-known property of a tree decomposition.

Lemma 2.1 (Lemma 12.3.1 in [14])**.** *Let $(\mathcal{T}, \beta)$ be a tree decomposition for a graph G. Then for every $t \in V(\mathcal{T})$ the pair $(\gamma(t), V(G) \setminus \alpha(t))$ is a separation of G; note that $\gamma(t) \cap (V(G) \setminus \alpha(t)) = \sigma(t)$.*

A *path decomposition* is a tree decomposition where $\mathcal{T}$ is a path, rooted at one of its endpoints.

(MaxEDP-LP)

$$\max \quad \sum_{i=1}^{k} x_i$$

$$\text{s.t.} \quad \sum_{p \in \mathcal{P}(s_i,t_i)} f(p) = x_i \leq 1, \quad i = 1, \ldots, k$$

$$\sum_{p:\, e \in p} f(p) \leq \text{cap}(e), \qquad e \in E(G)$$

$$f(p) \geq 0, \qquad p \in \mathcal{P} \ .$$

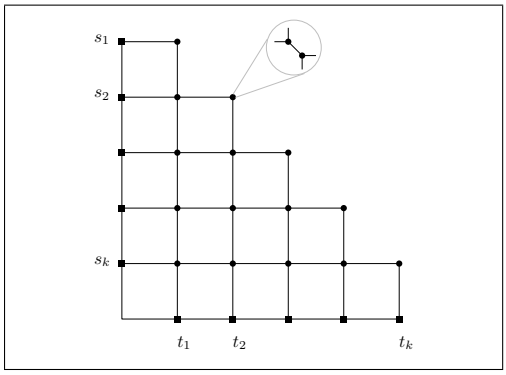

Figure 2: The multi-commodity flow relaxation for MaxEDP. The instance on the right is the $\Omega(\sqrt{n})$ integrality gap example for MaxEDP with unit edge capacities [16]. Any integral routing routes at most one pair whereas there is a multi-commodity flow that sends 1/2 units of flow for each pair (s_i, t_i) along the canonical path from s_i to t_i in the grid.

The width of a tree or path decomposition $(\mathcal{T}, \beta)$ is defined as $\max_t |\beta(t)| - 1$. To ease the notation, we will always consider decompositions of width *less* than r, for some integer r, so that every bag is of size at most r.

Problem definitions. The input to MaxEDP is an undirected graph G with edge capacities $\text{cap}(e) \in \mathbb{Z}_+$ and a collection $\mathcal{M} = \{(s_1, t_1), \ldots, (s_k, t_k)\}$ of vertex pairs. A *routing* for a subset $\mathcal{M}' \subseteq \mathcal{M}$ is a collection $\mathcal{P}$ of paths in G such that, for each pair $(s_i, t_i) \in \mathcal{M}'$, $\mathcal{P}$ contains a path connecting s_i to t_i. The routing is *feasible* if every edge e is in at most $\text{cap}(e)$ paths. In the Maximum Edge Disjoint Paths problem (MaxEDP), the goal is to maximize the number of pairs that can be feasibly routed. The Maximum Node Disjoint Paths problem (MaxNDP) is the node-capacitated variant of MaxEDP in which each node v has a capacity $\text{cap}(v)$ and in a feasible routing each node appears in at most $\text{cap}(v)$ paths.

We refer to the vertices participating in the pairs $\mathcal{M}$ as *terminals*. It is convenient to assume that $\mathcal{M}$ form a matching on the terminals; this can be ensured by making several copies of a terminal and attaching them as leaves.

Multicommodity flow relaxation. We use the following standard multicommodity flow relaxation for MaxEDP (there is an analogous relaxation for MaxNDP). We use $\mathcal{P}(u, v)$ to denote the set of all paths in G from u to v, for each pair (u, v) of nodes. Since the pairs $\mathcal{M}$ form a matching, the sets $\mathcal{P}(s_i, t_i)$ are pairwise disjoint. Let $\mathcal{P} = \bigcup_{i=1}^{k} \mathcal{P}(s_i, t_i)$. The LP has a variable $f(p)$ for each path $p \in \mathcal{P}$ representing the amount of flow on p. For each pair $(s_i, t_i) \in \mathcal{M}$, the LP has a variable x_i denoting the total amount of flow routed for the pair (in the corresponding IP, x_i denotes whether the pair is routed or not). The LP imposes the constraint that there is a flow from s_i to t_i of value x_i. Additionally, the LP has capacity constraints that ensure that the total amount of flow on paths using a given edge (resp. node for MaxNDP) is at the capacity of the edge (resp. node).

It is well-known that the relaxation MaxEDP-LP can be solved in polynomial time, since there is an efficient separation oracle for the dual (alternatively, one can write a compact relaxation). We use $(f, \mathbf{x})$ to denote a feasible solution to MaxEDP-LP for an instance $(G, \mathcal{M})$ of MaxEDP. For each terminal v, we also use $x(v)$ to denote the total amount of flow routed for v and we refer to $x(v)$ as the *marginal value* of v in the multi-commodity flow f.

3 Algorithm for MaxEDP in Bounded Treewidth Graphs

We give a polynomial time algorithm for MaxEDP that achieves an $O(r^3)$ approximation for graphs with treewidth less than r. Our algorithm builds on the work of Chekuri et al. [10], and it improves their approximation guarantee from $O(r \cdot 3^r)$ to $O(r^3)$. We use the following routing argument as a building block.

Proposition 3.1 (Proposition 3.4 in [10]). *Let $(G, \mathcal{M})$ be an instance of* MaxEDP *and let $(f, \mathbf{x})$ be a feasible fractional solution for the instance. If there is a second flow that routes at least $x(v)/\alpha$ units of flow*

for each v to some set $S \subseteq V(G)$, where $\alpha \geq 1$ then there is an integral routing of at least $\frac{|f|}{36\alpha|S|}$ pairs.

Our starting point is a tree decomposition $(\mathcal{T}, \beta)$ for G of width less than r and a fractional solution $(f, \mathbf{x})$ to the multicommodity flow relaxation for MaxEDP given in Section 2, that is, the flow f routes $\mathbf{x}(v)$ units of flow for each vertex $v \in V$. We let $|f|$ denote the total amount of flow routed by f, i.e., $|f| = \frac{1}{2}\sum_{v \in V} \mathbf{x}(v)$.

The following definitions play a key role in our algorithm.

Definition 3.2 (Safe node)**.** *A node $t \in V(\mathcal{T})$ is* safe *with respect to $(f, \mathbf{x})$ if there is a second multicommodity flow g in $G(t)$ such that g satisfies the edge capacities of $G(t)$ and, for each vertex $z \in \gamma(t)$, g routes at least $\frac{1}{4r} \cdot \mathbf{x}(z)$ units of flow from z to the adhesion $\sigma(t)$. The node t is* unsafe *if it is not safe.*

Definition 3.3 (Good node)**.** *A node $t \in V(\mathcal{T})$ is* good *with respect to $(f, \mathbf{x})$ if every flow path in the support of f that has an endpoint in $\gamma(t)$ also intersects $\sigma(t)$; in other words, there does not exist a flow path that is completely contained in $G[\alpha(t)]$. A node is* bad *if it is not good.*

Remark 3.4. *If a node t is good then it is also safe, as shown by the following multicommodity flow g in $G(t)$. For each path p in the support of f that originates in $\gamma(t)$, let p' be the smallest prefix of p that ends at a vertex of $\sigma(t)$ (since p intersects $\sigma(t)$, there is such a prefix); we set $g(p') = f(p)$. The resulting flow g is a feasible multicommodity flow in $G(t)$ that routes $\mathbf{x}(z)$ units of flow from z to $\sigma(t)$ for each vertex $z \in \gamma(t)$. Therefore, t is safe.*

Our approach is an inductive argument based on the maximum size of a parent adhesion that is bad or unsafe. More precisely, we prove the following:

Theorem 3.5. *Let $(G, \mathcal{M})$ be an instance of* MaxEDP *and let $(f, \mathbf{x})$ be a fractional solution for $(G, \mathcal{M})$, where f is a feasible multicommodity flow in G for $\mathcal{M}$ with marginals $\mathbf{x}$. Let $(\mathcal{T}, \beta)$ be a tree decomposition for G of width less than r. Let ℓ_1 be the maximum size of a parent adhesion of an unsafe node, and let ℓ_2 be the maximum size of a parent adhesion of a bad node. There is a polynomial time algorithm that constructs an integral routing of size at least $\frac{1}{144r^3} \cdot \left(1 - \frac{1}{r}\right)^{\ell_1 + \ell_2} \cdot |f|$.*

We start with a bit of preprocessing. If $|f| = 0$, then we return an empty routing. Otherwise, the root node of $\mathcal{T}$ is always unsafe and bad, and the integers ℓ_1 and ℓ_2 are well-defined. By considering every connected component of G independently (with inherited tree decomposition from $(\mathcal{T}, \beta)$), we assume that G is connected; note that in this step all safe or good adhesions remain safe or good for every connected component. Furthermore, we delete from $(\mathcal{T}, \beta)$ all nodes with empty bags; note that the connectivity of G ensures that the nodes with non-empty bags induce a connected subtree of $\mathcal{T}$. In this step, the root of $\mathcal{T}$ may have moved to a different node (the topmost node with non-empty bag), but the parent-children relation in the tree remains unchanged.

Once G is connected and no bag is empty, the only empty parent adhesion is the one for the root node. We prove Theorem 3.5 by induction on $\ell_1 + \ell_2 + |V(G)|$.

Base case. In the base case, we assume that $\ell_1 = \ell_2 = 0$. Since every parent adhesion of a non-root node is non-empty, that implies that the only bad node is the root t_0, that is, every flow path in f passes through $\beta(t_0)$, which is of size at most r. By applying Proposition 3.1 with $S = \beta(t_0)$ and $\alpha = 1$, we construct an integral routing of size at least $\frac{1}{36r}|f| \geq \frac{1}{144r^3}|f|$.

In the inductive step, we consider two cases, depending on whether $0 \leq \ell_1 < \ell_2$ or $0 < \ell_1 = \ell_2$.

Inductive step when $0 \leq \ell_1 < \ell_2$. Let $\{t_1, t_2, \ldots, t_p\}$ be the topmost bad nodes of $\mathcal{T}$ with parent adhesions of size ℓ_2, that is, it is a minimal set of such bad nodes such that for every bad node t with parent adhesion of size ℓ_2, there exists an $i \in \{1, \ldots, p\}$ with $t \preceq t_i$. For $1 \leq i \leq p$, let f_i^{inside} be the subflow of f consisting of all paths that are completely contained in $G[\alpha(t_i)]$. Furthermore, since $\ell_1 < \ell_2$, the node t_i is safe; let g_i be the corresponding flow, i.e., a flow that routes $\frac{1}{4r}x(v)$ from every $v \in \gamma(t_i)$ to $\sigma(t_i)$ in $G(t_i)$. By applying Proposition 3.1, there is an integral routing $\mathcal{P}_i$ in $G(t_i)$ that routes at least $\frac{1}{144r^2}|f_i^{\text{inside}}|$ pairs. Since the subgraphs $\{G(t_i) \colon 1 \leq i \leq p\}$ are edge-disjoint, we get an integral routing $\mathcal{P} \coloneqq \bigcup_i \mathcal{P}_i$ of size at least $\frac{1}{144r^2}\sum_{i=1}^{p}|f_i^{\text{inside}}|$.

If $\sum_{i=1}^{p}|f_i^{\text{inside}}| > \frac{1}{r}|f|$, then we can return the routing $\mathcal{P}$ as the desired solution. Otherwise, we drop the flows f_i^{inside}, that is, consider a flow $f' \coloneqq f - \sum_{i=1}^{p} f_i^{\text{inside}}$. Clearly, $|f'| \geq (1 - \frac{1}{r})|f|$. Furthermore, by definition of f_i^{inside}, every node t_i is good with respect to f'. Since deleting a flow path cannot turn a good

node into a bad one nor a safe node into an unsafe one, and all descendants of a good node are also good, we infer that every unsafe node with respect to f' has parent adhesion of size at most ℓ_1, while every bad node with respect to f' has parent adhesion of size *less* than ℓ_2. Consequently, by induction hypothesis we obtain an integral routing of size at least

$$\frac{1}{144r^3}\left(1-\frac{1}{r}\right)^{\ell_1+\ell_2-1}|f'| \geq \frac{1}{144r^3}\left(1-\frac{1}{r}\right)^{\ell_1+\ell_2}|f|.$$

Inductive step when $0 < \ell_1 = \ell_2$. In this case, we pick a node t° to be the lowest node of $\mathcal{T}$ that is unsafe and has parent adhesion of size ℓ_1. By the definition of an unsafe node and Menger's theorem, there exists a set $U \subseteq \alpha(t^\circ)$ such that $\mathrm{cap}(\delta(U)) < \frac{1}{4r}\mathbf{x}(U)$. With a bit more care, we can extract a set U with one more property:

Lemma 3.6. *In polynomial time we can find a set $U \subseteq \alpha(t^\circ)$ for which*

1. $\mathrm{cap}(\delta(U)) < \frac{1}{4r}\mathbf{x}(U)$;
2. *for every non-root node t, if $\sigma(t) \subseteq U$, then $\gamma(t) \subseteq U$.*

Proof. Consider an auxiliary graph G', obtained from $G[\gamma(t^\circ)]$ by adding a super-source s^*, linked for every $v \in \gamma(t^\circ)$ by an arc (s^*, v) of capacity $\frac{1}{4r}\mathbf{x}(v)$, and a super-sink t^*, linked for every $v \in \sigma(t^\circ)$ by an arc (v, t^*) of infinite capacity. Let U be such a set that $\delta(U \cup \{s^*\})$ is a minimum s^*-t^* cut in this graph. Clearly, since U is unsafe, $\mathrm{cap}(\delta_{G'}(U \cup \{s^*\})) < \frac{1}{4r}\mathbf{x}(\gamma(t^\circ)) = \mathrm{cap}(\delta_{G'}(s^*))$, so $U \neq \emptyset$. Also, $U \subseteq \alpha(t^\circ)$, as each node in $\sigma(t^\circ)$ is connected to t^* with an infinite-capacity arc.

We claim that U satisfies the desired properties. The first property is immediate:

$$\mathrm{cap}(\delta_G(U)) = \mathrm{cap}(\delta_{G'}(U \cup \{s^*\})) - \frac{1}{4r}\mathbf{x}(\gamma(t^\circ) \setminus U) < \frac{1}{4r}\left(\mathbf{x}(\gamma(t^\circ)) - \mathbf{x}(\gamma(t^\circ) \setminus U)\right) = \frac{1}{4r}\mathbf{x}(U).$$

For the second property, pick a non root node t with $\sigma(t) \subseteq U$. Since $\sigma(t) \subseteq U \subseteq \alpha(t^\circ)$, we have $t \preceq t^\circ$, $t \neq t^\circ$, and $\gamma(t) \subseteq \alpha(t^\circ)$. Let $U' \coloneqq U \cup \gamma(t)$. By Lemma 2.1, $\delta_G(U') \subseteq \delta_G(U)$, and hence $\delta_{G'}(U' \cup \{s^*\}) \subseteq \delta_{G'}(U \cup \{s^*\})$. However, since $\delta_{G'}(U \cup \{s^*\})$ is a minimum cut, we have actually $\delta_G(U') = \delta_G(U)$. Since G is connected, this implies that $U = U'$, and thus $\gamma(t) \subseteq U$. As the choice of t was arbitrary, U satisfies the second property. □

Using the cut U, we split the graph G and the flow f into two pieces as follows. Let $G_1 = G[U]$ and $G_2 = G - U$. Let f_i be the restriction of f to G_i, i.e., the flow consisting of only flow paths that are contained in G_i. Let $\mathbf{x}_i$ be the marginals of f_i and let $\mathcal{M}_i$ be the subset of $\mathcal{M}$ consisting of all pairs (s,t) such that $\{s,t\} \subseteq V(G_i)$; note that $x_i(s) = x_i(t)$ for each pair $(s,t) \in \mathcal{M}_i$ and thus $(f_i, \mathbf{x}_i)$ is a fractional routing for the instance $(G_i, \mathcal{M}_i)$. Let $(\mathcal{T}, \beta_1)$ and $(\mathcal{T}, \beta_2)$ be the restriction of $(\mathcal{T}, \beta)$ to the vertices of G_1 and G_2, respectively; we define mappings σ_i, γ_i, and α_i naturally. In what follows, we consider separately two instances $\mathcal{I}_i \coloneqq \langle (G_i, \mathcal{M}_i), (f_i, \mathbf{x}_i), (\mathcal{T}, \beta_i)\rangle$ for $i = 1, 2$.

An important observation is the following:

Lemma 3.7. *Every node $t \in V(\mathcal{T})$ that is good in the original instance (i.e., as a node of $\mathcal{T}$, with respect to $(f, \mathbf{x})$) is also good in $\mathcal{I}_i$ with respect to $(f_i, \mathbf{x}_i)$.*

Proof. Note that every flow path in f_i is also present in f, and therefore intersects the parent adhesion of f if t is a good node in the original instance. □

Consequently, every node $t \in V(\mathcal{T})$ with $|\sigma(t)| > \ell_2$ is good in the instance $\mathcal{I}_i$, and the maximum size of a parent adhesion of a bad node in instance $\mathcal{I}_i$ is at most ℓ_2. Hence, both $\mathcal{I}_1$ and $\mathcal{I}_2$ satisfy the assumptions of Theorem 3.5 with not larger values of ℓ_1 and ℓ_2. Furthermore, note that $|V(G_i)| < |V(G)|$ for $i = 1, 2$.

For $\mathcal{I}_2$, the above reasoning allows us to simply just apply inductive step, obtaining an integral routing $\mathcal{P}_2$ of size at least

$$|\mathcal{P}_2| \geq \frac{1}{144r^3}\left(1-\frac{1}{r}\right)^{\ell_1+\ell_2} \cdot |f_2| \ . \tag{1}$$

For $\mathcal{I}_1$, we are going to obtain a larger routing via an inductive step with better bounds.

Lemma 3.8. *The size of the largest parent adhesion of an unsafe note in $\mathcal{I}_1$ is* less *than ℓ_1.*

Proof. Assume the contrary, let $t \in V(\mathcal{T})$ be an unsafe adhesion with $|\sigma_1(t)| \geq \ell_1$. If $|\sigma(t)| > \ell_1$, then t is good in the original instance, and by Lemma 3.7 it remains good in $\mathcal{I}_1$. Consequently, $|\sigma(t)| = |\sigma_1(t)| = \ell_1$; in particular, $\sigma(t) = \sigma_1(t) \subseteq U$.

By Lemma 3.6, property 2, we have $\gamma(t) \subseteq U$. Consequently, t is safe in the original instance if and only if it is safe in $\mathcal{I}_1$. Since $t \preceq t^\circ$, $t \neq t^\circ$, but $|\sigma(t)| = \ell_2$, by the choice of t° it holds that t is safe in the original instance, a contradiction. □

Lemma 3.8 allows us to apply the inductive step to $\mathcal{I}_1$ and obtain an integral routing $\mathcal{P}_1$ of size at least

$$|\mathcal{P}_1| \geq \frac{1}{144r^3}\left(1-\frac{1}{r}\right)^{\ell_1-1+\ell_2} \cdot |f_1| \ . \tag{2}$$

Let us now estimate the amount of flow lost by the separation into $\mathcal{I}_1$ and $\mathcal{I}_2$, i.e., $g = f - f_1 - f_2$. As every flow path in g passes through $\delta(U)$, we have $|g| \leq \mathrm{cap}(\delta(U)) < \frac{1}{4r}\mathbf{x}(U)$. Since $|f_1| + |g| \geq \frac{1}{2}\mathbf{x}(U)$ (no flow path in f_2 originates in U), we have that $|g| \leq \frac{1}{4r} \cdot 2 \cdot (|f_1| + |g|)$. Hence,

$$|g| \leq \frac{1}{2r} \cdot \left(1-\frac{1}{2r}\right)^{-1} |f_1| \leq \frac{1}{r}|f_1| \ . \tag{3}$$

By putting up together (1), (2), and (3), we obtain that

$$\begin{aligned} |\mathcal{P}_1| + |\mathcal{P}_2| &\geq \frac{1}{144r^3}\left(1-\frac{1}{r}\right)^{\ell_1+\ell_2}\left(|f_2| + \left(1-\frac{1}{r}\right)^{-1}|f_1|\right) \\ &\geq \frac{1}{144r^3}\left(1-\frac{1}{r}\right)^{\ell_1+\ell_2}(|f_2| + |f_1| + |g|) = \frac{1}{144r^3}\left(1-\frac{1}{r}\right)^{\ell_1+\ell_2}|f| \ . \end{aligned}$$

This concludes the proof of Theorem 3.5. Since $\ell_1, \ell_2 \leq r$, while $(1-\frac{1}{r})^{2r} = \Omega(1)$, Theorem 3.5 immediately implies the promised $\mathcal{O}(r^3)$-approximation algorithm.

Remark 3.9. *We conclude with observing that the improved approximation ratio of $\mathcal{O}(r^3)$ directly translates to the more general setting of k-sums of graph from some minor closed family $\mathcal{G}$, as discussed in [10]. That is, if we are able to α-approximate* MaxEDP *with congestion β in graphs from $\mathcal{G}$, we can have $\mathcal{O}(\alpha r^5)$-approximation algorithm with congestion $(\beta+3)$ in graphs admitting a tree decomposition of maximum adhesion size at most r, and the torso of every bag being from the class $\mathcal{G}$.*

To see this, observe that the only place when our algorithm uses that the bags *are of bounded size (as opposed to* adhesions*) is the base case, where all flow paths pass through the bag $\beta(t_0)$ of the root node t_0. However, in this case we can proceed exactly as Chekuri et al. [10]: using the flow paths, move the terminals to $\beta(t_0)$, replace connected components of $G - \beta(t_0)$ with their $(r^2, 2)$-sparsifiers, and apply the algorithm for the class $\mathcal{G}$. In addition to the $\mathcal{O}(r^3)$ approximation factor of our algorithm, the application of the algorithm for $\mathcal{G}$ incurs an approximation ratio of α and congestion of β, the use of sparsifiers adds a factor of r^2 to the approximation ratio and an additive constant $+1$ to the congestion, while the terminal move adds an additional amount of 2 to the final congestion.*

4 Algorithm for MaxNDP in Bounded Pathwidth Graphs

In this section we develop an $\mathcal{O}(r^3)$-approximation algorithm for MaxNDP in graphs of *pathwidth* less than r. We follow the outline of the MaxEDP algorithm from the previous section, with few essential changes.

Most importantly, we can no longer use Proposition 3.1, as it refers to edge disjoint paths, and the proof of its main ingredient by Chekuri et al. [6] relies on a clustering technique that stops to work for node disjoint paths. We fix this issue by providing in Sect. 4.1 a node-disjoint variant of Proposition 3.1, using the more involved clustering approach of Chekuri et al. [7].

Then, in Sect. 4.2 we revisit step-by-step the arguments for MaxEDP, pointing out remaining differences. We remark that the use of pathwidth instead of treewidth is only essential in the inductive step for the case $\ell_1 < \ell_2$: if we follow the argument for MaxEDP for bounded-treewidth graphs, the graphs $G(t_i)$ may not be node disjoint (but they are edge disjoint), breaking the argument. Note that for bounded pathwidth graphs, there is only one such graph considered, and the issue is nonexistent.

4.1 Routing to a small adhesion in a node-disjoint setting

In this section we prove the following statement.

Proposition 4.1. *Let $(G, \mathcal{M})$ be an instance of* MaxNDP *and let $(f, \mathbf{x})$ be a feasible fractional solution for the instance. Suppose that there is also a second (feasible, i.e., respecting node capacities) flow that routes at least $x(v)/\alpha$ units of flow for each v to some set $S \subseteq V$, where $\alpha \geq 1$. Then there is an integral routing of $\Omega(|f|/(\alpha|S|))$ pairs.*

Without loss of generality, we may assume that the terminals of $\mathcal{M}$ are pairwise distinct and of degree and capacity one: we can always move a terminal from a vertex t to a newly-created degree-1 capacity-1 neighbour of t.

Let g be the second flow mentioned in the statement. In what follows, we modify and simplify the flows f and g in a number of steps. We denote by $f_1, f_2, \ldots$ and $g_1, g_2, \ldots$ flows after subsequent modification steps; for the flow f_i, by $\mathbf{x}_i$ we denote its marginals.

Symmetrizing the flow g. In the first step, we construct flows f_1 and g_1 with the following property: for every terminal pair $(s,t) \in \mathcal{M}$, for every $v \in S$, g_1 sends the same amount of flow from s to v as from t to v. To obtain this goal, we first take the flow $g/3$, and then for every $(s,t) \in \mathcal{M}$ redirect the flow originating at s to first go along the commodity for the pair (s,t) in flow $f/(3\alpha)$ to the vertex t, and then go to S in exactly the same manner as the flow originating at t does. It is easy to see that g_1 consists of three feasible flows scaled down by at least $1/3$, thus it is feasible. Finally, we set $f_1 \coloneqq f/3$, so that g_1 again sends $\mathbf{x}_1(v)/\alpha$ flow from every vertex v to S. Note that $|f_1| = |f|/3$.

Restricting to single vertex of S. To construct flows f_2 and g_2, pick a vertex $u \in S$ that receives the most flow in g_1. Take g_2 to be the flow g_1, restricted only to flow paths ending in u. Then, restrict f_1 to obtain f_2 as follows: for every terminal pair $(s,t) \in \mathcal{M}$, reduce the amount of flow from s to t to α times the total amount of flow sent from s to u by g_2; note that, by the previous step, it is also equal α times the total amount of flow sent from t to u by g_2. By this step, we maintain the invariant that g_2 sends $\mathbf{x}_2(v)/\alpha$ flow from every $v \in V(G)$, and we have $|f_2| \geq |f_1|/|S| \geq |f|/(3|S|)$.

Rounding to a half-integral flow. In the next step, we essentially repeat the integral rounding procedure by Chekuri et al. [6, Section 3.2]. We use the following operation as a basic step in the rounding.

Lemma 4.2 (Theorem 2.1 of [6])**.** *Let G be a directed graph with edge capacities. Given a flow h in G that goes from set $X \subseteq V(G)$ to a single vertex $u \in V(G)$, such that for every $v \in X$ the amount of flow originating in v is $\mathbf{z}(v)$, and a vertex $v_0 \in X$ such that $\mathbf{z}(v_0)$ is not an integer, one can in polynomial time compute a flow h' in G, sending $\mathbf{z}'(v)$ amount of flow from every $v \in X$ to u, such that $|h'| \geq |h|$, $\mathbf{z}'(v) = \mathbf{z}(v)$ for every $v \in X$ where $\mathbf{z}(v)$ is an integer, and $\mathbf{z}'(v_0) = \lceil \mathbf{z}(v_0) \rceil$.*

Since a standard reduction reduces flows in undirected node-capacitated graphs to directed edge-capacitated ones[3], Lemma 4.2 applies also to undirected graphs with integral node capacities.

Split g_2 into two flows h_s and h_t: for every terminal pair $(s,t) \in \mathcal{M}$, we put the flow originating in s into h_s, and the flow originating in t into h_t. We perform a sequence of modifications to the flows h_s and h_t, maintaining the invariant that the same amount of flow originates in s in h_s as in t in h_t for every $(s,t) \in \mathcal{M}$. Along the process, both h_s and h_t are feasible flows, but $h_s + h_t$ may not be.

In a single step, we pick a terminal pair $(s,t) \in \mathcal{M}$ such that the amount of flow in h_s originating in s is not integral (and stop if no such pair exists). We apply Lemma 4.2 separately to s in h_s and to t in h_t,

[3]Replace every edge with two infinite-capacity arcs in both directions, and then split every vertex into two vertices, connected by an edge of capacity equal to the capacity of the vertex, with all in-edges connected to the first copy, and all out-edges connected to the second copy.

obtaining flows h'_s and h'_t. Finally, if for some terminal pair (s',t'), the amount of flow originating in s' in h'_s and in t' in h'_t differ, we restrict one of the flows so that both route the same amount of flow (being the minimum of the flows routed by h'_s from s' and by h'_t from t').

Since the rounding algorithm of Lemma 4.2 never modifies a source that already has an integral flow, this procedure stops after at most $|\mathcal{M}|$ steps. Furthermore, if in one step the flow from s has been increased from z to $\lceil z \rceil$, the total loss of flow to other pairs is $2(\lceil z \rceil - z)$. Therefore, if h°_s and h°_t are the final integral flows, we have $|h^\circ_s| + |h^\circ_t| \geq (|h_s| + |h_t|)/2 = |g_2|/2 = |f_2|/\alpha \geq |f|/(3\alpha|S|)$. Finally, we define $g_3 := (h^\circ_s + h^\circ_t)/2$; note that g_3 is a feasible flow, since both h°_s and h°_t are.

Clustering a node-flow-linked set. Note that for every $(s,t) \in \mathcal{M}$, the flow g_3 routes either 0 or $1/2$ flow from both s and t to u. Let $\mathcal{M}'$ be the set of pairs for which the flow is $1/2$, and let X' be the set of terminals in $\mathcal{M}'$. Note that $|\mathcal{M}'| = |g_3|/2 \geq |f|/(6\alpha|S|)$.

Observe that the set X' is $\frac{1}{4}$-node-flow-linked: using g_3, we can find a multicommodity flow that for every $(a,b) \in X' \times X'$ routes $\frac{1}{4|X|}$ amount of flow from a to b, by routing it along g_3 to u and along reversed g_3 from u. This allows us to apply the following clustering result.

Lemma 4.3 (Lemma 2.7 of [7])**.** *If X is α-node-flow-linked in a graph G with unit node capacities, then for any $h \geq 2$ there exists a forest F in G of maximum degree $\mathcal{O}(\frac{1}{\alpha}\log h)$ such that every tree in F spans at least h nodes from X.*

Since we can assume that no capacity in G exceeds $|\mathcal{M}|$, we can replace every vertex v of capacity $\mathrm{cap}(v)$ with its $\mathrm{cap}(v)$ copies. To such unweighted graph G' we apply Lemma 4.3 for X', $\alpha = 1/4$ and $h = 3$, obtaining a forest F'; recall that the terminals X' are of capacity 1, thus they are kept unmodified in G'. By standard argument we split the forest F' into node-disjoint trees $T'_1, T'_2, \ldots, T'_p$, such that every tree T'_i contains at least three, and at most $d = \mathcal{O}(1)$ terminals of X'. By projecting the trees T'_i back onto G, we obtain a sequence of trees $T_1, T_2, \ldots, T_p$, such that every vertex $v \in V(G)$ is present in at most $\mathrm{cap}(v)$ trees T_i. Furthermore, since terminals are of capacity one, every terminal belongs to at most one tree, and every tree T_i contains at least three and at most d terminals.

In a greedy fashion, we chose a set $\mathcal{M}'' \subseteq \mathcal{M}'$ of size at least $|\mathcal{M}'|/d^2$, such that for every tree T_i, at most one terminal pair of $\mathcal{M}''$ has at least one terminal in T_i. A pair $(s,t) \in \mathcal{M}''$ is *local* if both s and t lie in the same tree T_i, and *distant* otherwise. If at least half of the pairs of $\mathcal{M}''$ are local, we can route them along trees T_i, obtaining a desired routing of size at least $|\mathcal{M}''|/2 \geq |\mathcal{M}'|/(2d^2) = \Omega(|f|/(\alpha|S|))$ and terminate the algorithm. Otherwise, we obtain a flow g_4 as follows: for every terminal t in a distant pair in $\mathcal{M}''$, we take the tree T_i it lies on, route $3/5$ amount of flow along T_i equidistributed to three arbitrarily chosen terminals t^1, t^2, t^3 on T_i from $\mathcal{M}'$ (i.e., every terminal t^j receives $1/5$ amount of flow), and then route the flow along the flow $\frac{2}{5}g_3$ to u. Since every tree T_i routes $3/5$ amount of flow, and g_3 is a feasible flow, the flow g_4 is a feasible flow that routes $3/5$ amount of flow from every terminal of $\mathcal{M}''$ to u. Furthermore, since at least half terminal pairs in $\mathcal{M}''$ is distant, we have $|g_4| \geq \frac{1}{2} \cdot 2|\mathcal{M}''| = \Omega(|f|/(\alpha|S|))$.

Final rounding of the flow. Let X'' be the set of all terminals of $\mathcal{M}''$. Since the flow g_4 routes *more than* $1/2$ amount of flow for every terminal in X'', we can conclude with simple rounding the flow g_4 in the same manner as it is done by Chekuri et al. [6, Section 3]. Construct an auxiliary graph G' by adding a super-source s^* of infinite capacity, adjacent to all terminals of $\mathcal{M}''$. Extend g_4 in the natural manner, by routing every flow path first from s^* to an appropriate terminal. The extended flow g_4 is now a single souce single sink flow from s^* to u in a graph with integer capacities, thus there exists an integral flow g_5 of no smaller size:

$$|g_5| \geq |g_4| = \frac{3}{5}|X''| = \frac{6}{5}|\mathcal{M}''|.$$

Hence, for at least $1/5$ of the pairs $(s,t) \in \mathcal{M}''$, the flow g_5 routes a single unit of flow both from s and from t to u. By combining these paths into a single path from s to t, we obtain an integral routing of size at least $\frac{1}{5}|\mathcal{M}''| = \Omega(|f|/(\alpha|S|))$. This finishes the proof of Proposition 4.1.

4.2 Details of the algorithm

Equipped with Proposition 4.1, we can now proceed to the description of the approximation algorithm. Assume we are given an MaxNDP instance $(G, \mathcal{M})$ and a path decomposition $(\mathcal{T}, \beta)$ of G of width less than r; recall that $\mathcal{T}$ rooted in one of its endpoints. Let $(f, \mathbf{x})$ be a fractional solution to the multicommodity flow relaxation for MaxNDP, as in Section 2.

The definitions of safe and good node are analogous, and we follow the same induction scheme.

Definition 4.4 (Safe node)**.** *A node $t \in V(\mathcal{T})$ is* safe *with respect to $(f, \mathbf{x})$ if there is a second multicommodity flow g in $G(t)$ such that g satisfies the node capacities of $G(t)$ and, for each vertex $z \in \gamma(t)$, g routes at least $\frac{1}{4r} \cdot \mathbf{x}(z)$ units of flow from z to the adhesion $\sigma(t)$. The node t is* unsafe *if it is not safe.*

Definition 4.5 (Good node)**.** *A node $t \in V(\mathcal{T})$ is* good *with respect to $(f, \mathbf{x})$ if every flow path in the support of f that has an endpoint in $\gamma(t)$ also intersects $\sigma(t)$; in other words, there does not exist a flow path that is completely contained in $G[\alpha(t)]$. A node is* bad *if it is not good.*

Theorem 4.6. *Let $(G, \mathcal{M})$ be an instance of* MaxNDP *and let $(f, \mathbf{x})$ be a fractional solution for the instance, where f is a feasible multicommodity flow in G for the pairs $\mathcal{M}$ with marginals $\mathbf{x}$. Let $(\mathcal{T}, \beta)$ be a path decomposition for G of width less than r. Let ℓ_1 be the maximum size of a parent adhesion of an unsafe node, and let ℓ_2 be the maximum size of a parent adhesion of a bad node. There is a constant c and a polynomial time algorithm that constructs an integral routing of size at least*

$$\frac{1}{cr^3} \cdot \left(1 - \frac{1}{r}\right)^{\ell_1 + \ell_2} \cdot |f| \ .$$

Again as in the case of MaxEDP, we can assume that the considered graph G is connected and that no bag is empty, and thus the only empty adhesion is the parent adhesion of the root.

Base case. In the base case $\ell_1 = \ell_2 = 0$ nothing changes as compared to MaxEDP: all flow paths pass through the root bag, and Proposition 4.1 allows us to route integrally $\Omega(|f|/r)$ paths.

Inductive step when $0 \leq \ell_1 < \ell_2$. Since we are considering now a path decomposition (as opposed to tree decomposition in the previous section), there exists a single topmost bad node t° with parent adhesion of size ℓ_2. Define f^{inside} to be the subflow of f consisting of all flow paths completely contained in $G[\alpha(t^\circ)]$. Since $\ell_1 < \ell_2$, the node t° is safe, and the flow witnessing it together with Proposition 4.1 allows us to integrally route $\Omega(|f^{\text{inside}}|/r^2)$ terminal pairs. If $|f^{\text{inside}}| > |f|/r$, then we are done. Otherwise, we drop the flow f^{inside} from f, making t° and all its descendants good (thus decreasing ℓ_2 in the constructed instance), while losing only $1/r$ fraction of the flow f, and pass the instance to an inductive step.

Inductive step when $0 < \ell_1 = \ell_2$. Here again we take t° to be the lowest node of $\mathcal{T}$ that is unsafe and has parent adhesion of size ℓ_1. By the definition of an unsafe node and Menger's theorem, there exists a set $U \subseteq \alpha(t^\circ)$ such that $\mathrm{cap}(N(U)) < \frac{1}{4r}\mathbf{x}(U)$. Using the same argument as in the proof of Lemma 3.6, we can ensure property 2, that is that if U contains an adhesion $\sigma(t)$, it contains as well the entire set $\gamma(t)$.

As in the case of MaxEDP, we split into instances $\mathcal{I}_1$ and $\mathcal{I}_2$ by taking $G_1 = G[U]$ and $G_2 = G - N[U]$, with inherited tree decompositions from $(\mathcal{T}, \beta)$. Since all nodes with parent adhesions of size larger than $\ell_1 = \ell_2$ are good, there are also good in instances $\mathcal{I}_i$ (i.e., Lemma 3.7 holds here as well) and we can again apply the inductive step to every connected component of the instance $\mathcal{I}_2$ with the same values of ℓ_1 and ℓ_2, obtaining a routing $\mathcal{P}_2$ of size as in (1) (with 144 replaced by a constant c).

We analyse the instance $\mathcal{I}_1$, without breaking it first into connected components. That is, we argue that in $\mathcal{I}_1$ the value of ℓ_1 dropped, that is, all nodes t satisfying $|\sigma(t)| = |\sigma_1(t)| = \ell_1$ are safe; note that they will remain safe once we consider every connected component separatedly. However, this fact follows from property 2 of the set U (Lemma 3.6): if for some node t we have $|\sigma(t)| = |\sigma_1(t)|$, it follows that $\sigma(t) \subseteq U$ hence $\gamma(t) \subseteq U$ and the notion of safeness for t is the same in $\mathcal{I}_1$ and in the original instance. However, $\sigma(t) \subseteq U \subseteq \alpha(t^\circ)$ implies $t \preceq t^\circ$ and $t \neq t^\circ$, hence t is safe in the original instance.

Consequently, an application of inductive step for every connected component of $\mathcal{I}_1$ uses strictly smaller value of ℓ_1, and we obtain an integral routing $\mathcal{P}_1$ in $\mathcal{I}_1$ of size as in (2) (again with 144 replaced by a constant c). The remainder of the analysis from the previous section does not change, concluding the proof of Theorem 4.6.□

5 Well-linked decomposition

In this section, we show that the argument given in Section 3 gives the following well-linked decomposition for edge-capacitated graphs. We first give some preliminary definitions.

Multi-commodity flows. We represent a multi-commodity flow instance as a demand vector $\mathbf{d}$ that assigns a demand $d(u,v) \in \mathbb{R}_+$ to each ordered pair (u,v) of vertices of G. A *product* multi-commodity flow instance satisfies $d(u,v) = w(u)w(v)$ for each pair (u,v), where $w : V \to \mathbb{R}_+$ is a weight functions on the vertices of G. We say that $\mathbf{d}$ is *routable* if there is a feasible multi-commodity flow in G that routes $d(u,v)$ units of flow from u to v for each pair (u,v).

We recall the following two quantities associated with a multi-commodity flow instance: the maximum concurrent flow and the sparsest cut. The *maximum concurrent flow* is the maximum value $\lambda \geq 0$ such that $\lambda\mathbf{d}$ is routable. The *sparsity* of a cut $S \subseteq V$ is the ratio $|\delta(S)|/\mathrm{dem}_{\mathbf{d}}(S)$, where $\mathrm{dem}_{\mathbf{d}}(S) = \sum_{u\in S, v\in V\setminus S} d(u,v)$ is the total demand separated by S. A *sparsest cut* is a cut with minimum sparsity. The minimum sparsity of any cut is an upper bound on the maximum concurrent flow, and the former could be strictly larger. The *flow-cut gap* is the worst-case ratio between the sparsest cut and the maximum concurrent flow. The flow-cut gap is $\mathcal{O}(\log k)$ in general graphs, where k is the number of commodities (that is, the number of non-zero demands $d(u,v)$) [20]. For product multi-commodity flows, the flow-cut gap is $\mathcal{O}(\log r)$ for graphs of treewidth r [9]. Moreover, for both of these results, there are polynomial time algorithms for computing a cut with sparsity at most $\alpha(G)\cdot\lambda$, where $\alpha(G)$ is the flow-cut gap value and λ is the maximum concurrent flow for $\mathbf{d}$ [20, 9].

Well-linked sets. Following [7], we work with two notions of well-linked sets, cut well-linked sets and flow well-linked sets. For convenience, we work with graphs with unit edge capacities; it is straightforward to extend the argument to arbitrary edge capacities.

Let $\pi : X \to \mathbb{R}_+$ be a weight function on a set $X \subseteq V(G)$ of vertices. The set X is *π-flow-well-linked* in G if there is a feasible multi-commodity flow that simultaneously routes $d(u,v) \coloneqq \pi(u)\pi(v)/\pi(X)$ units of flow from u to v for every pair (u,v) of nodes in X. The set X is *π-cut-well-linked* in G if $|\delta(S)| \geq \min\{\pi(S\cap X), \pi((V\setminus S)\cap X)\}$.

Chekuri et al. [7] gave the following well-linked decomposition theorem for general graphs that has found many applications. We only state the theorem for flow-well-linked instances, there is an analogous result for cut-well-linked instances.

Theorem 5.1 ([7])**.** *Let* OPT *be the value of a solution to* MaxEDP-LP *for a given instance $(G,\mathcal{M})$ of* MaxEDP *on a general graph G. Let $\alpha = \alpha(G) \geq 1$ be an upper bound on the worst case flow-cut gap for product multi-commodity flows in G. There is a partition of G into node-disjoint induced subgraphs $G_1, G_2, \ldots, G_q$ and weight functions $\pi_i : V(G_i) \to \mathbb{R}_+$ with the following properties. Let $\mathcal{M}_i$ be the induced pairs of $\mathcal{M}$ in G_i and let X_i be the endpoints of the pairs in $\mathcal{M}_i$. We have*

(a) $\pi_i(u) = \pi_i(v)$ for each pair $uv \in \mathcal{M}_i$.

(b) X_i is π_i-flow-well-linked in G_i.

(c) $\sum_{i=1}^{q} \pi_i(X_i) = \Omega(\mathrm{OPT}/(\alpha \log \mathrm{OPT})) = \Omega(\mathrm{OPT}/\log^2 k)$.

Moreover, such a partition is computable in polynomial time if there is a polynomial time algorithm for computing a node separator with sparsity at most $\alpha(G)$ times the maximum concurrent flow.

In the remainder of this section, we show that, for graphs with treewidth at most r, the argument in Sect. 3 implies a well-linked decomposition that loses a factor of $\mathcal{O}(r^3)$ instead of $\mathcal{O}(\log^2 k)$.

Theorem 5.2. *Let* OPT *be the value of a solution to* MaxEDP-LP *for a given instance $(G,\mathcal{M})$ of* MaxEDP *on a graph G of treewidth at most r. There is a partition of G into node-disjoint induced subgraphs $G_1, G_2, \ldots, G_q$ and weight functions $\pi_i : V(G_i) \to \mathbb{R}_+$ with the following properties. Let $\mathcal{M}_i$ be the induced pairs of $\mathcal{M}$ in G_i and let X_i be the endpoints of the pairs in $\mathcal{M}_i$. We have*

(a) $\pi_i(u) = \pi_i(v)$ for each pair $uv \in \mathcal{M}_i$.

(b) There is a feasible multi-commodity flow g_i in G_i from X_i to a single vertex z that routes $\pi_i(v)$ units of flow from v to z for each vertex $v \in X_i$. Thus X_i is π_i-flow-well-linked in G_i.

(c) $\sum_{i=1}^{q} \pi_i(X_i) = \Omega(\mathrm{OPT}/r^3)$.

Moreover, such a partition is computable in polynomial time if one is given a tree decomposition of G of width at most r.

In the remainder of this section, we show how to modify the argument in Section 3 to prove Theorem 5.2. We use the definitions and notation introduced in Section 3.

As before, our starting point is a fractional solution $(f, \mathbf{x})$ to MaxEDP-LP for the instance $(G, \mathcal{M})$. We prove the following theorem by induction on the maximum size of a parent adhesion that is bad or unsafe.

Theorem 5.3. *Let $(G, \mathcal{M})$ be an instance of* MaxEDP *and let $(f, \mathbf{x})$ be a fractional solution for the instance, where f is a feasible multicommodity flow in G for the pairs $\mathcal{M}$ with marginals $\mathbf{x}$. Let $(\mathcal{T}, \beta)$ be a tree decomposition for G of width less than r. Let ℓ_1 be the maximum size of a parent adhesion of an unsafe node, and let ℓ_2 be the maximum size of a parent adhesion of a bad node. There is a polynomial time algorithm that constructs a partition of G into node-disjoint induced subgraphs $G_1, G_2, \ldots, G_q$ and weight functions $\pi_i : V(G_i) \to \mathbb{R}_+$ with the following properties. Let $\mathcal{M}_i$ be the induced pairs of $\mathcal{M}$ in G_i and let X_i be the endpoints of the pairs in $\mathcal{M}_i$. We have*

(a) $\pi_i(u) = \pi_i(v)$ for each pair $uv \in \mathcal{M}_i$.

(b) There is a feasible multi-commodity flow g_i in G_i from X_i to a single vertex z that routes $\pi_i(v)$ units of flow from v to z for each vertex $v \in X_i$. Thus X_i is π_i-flow-well-linked in G_i.

(c) $\sum_{i=1}^{q} \pi_i(X_i) \geq \frac{1}{12r^3} \cdot \left(1 - \frac{1}{r}\right)^{\ell_1+\ell_2} \cdot |f|$.

In addition to the definitions of a good and safe node from Section 3, it is convenient to have the following definition.

Definition 5.4 (Nice flow)**.** *A multi-commodity flow f is* nice *for a node $t \in V(\mathcal{T})$ if there exists a single vertex $z \in \sigma(t)$ such that each flow path P in the support of f ends at z and $P - z$ is contained in $G[\alpha(t)]$.*

The proof of the following lemma follows from a result by Chekuri et al. [10, Proposition 3.4], and we include it for completeness.

Lemma 5.5. *Consider a node $t \in V(\mathcal{T})$. Let $\mathcal{M}$ be a collection of pairs with both endpoints in $\alpha(t)$ such that each vertex of $\alpha(t)$ appears in at most one pair. Let f be a multi-commodity flow in $G[\alpha(t)]$ with marginals $\mathbf{x}$ satisfying $x(u) = x(v)$ for each pair $(u, v) \in \mathcal{M}$. Suppose that there is a second multi-commodity flow g in $G(t)$ that routes at least $x(v)/c$ units of flow for each vertex $v \in V(\mathcal{M})$ to $\sigma(t)$, where $c \geq 1$. Then there is a multi-commodity flow h from $V(\mathcal{M})$ to z with the following properties: h is nice for t; for each pair $(u, v) \in \mathcal{M}$, h routes the same amount of flow for u and v; $|h| \geq \frac{1}{3c|\sigma(t)|} \cdot |f| \geq \frac{1}{3cr} \cdot |f|$.*

Proof. Note that we may assume that each flow path in the support of $G(t)$ is contained in $G[\alpha(t)]$ except for its end vertex. Let $z \in \sigma(t)$ be a vertex that receives the most g-flow, where the g-flow received by z is the total amount of flow on paths ending at z. Let g' be the sub-flow of g consisting of only the flow paths ending at z, and let $\mathbf{x}'$ be the marginals of g'.

Note that $x'(v) \leq x(v)/c$ for every $v \in V(\mathcal{M})$ and $|g'| \geq \frac{1}{c|\sigma(t)|} \sum_{v \in \alpha(t)} |f|$. However, there may be pairs $(u, v) \in \mathcal{M}$ for which $x'(u) \neq x'(v)$. We use the flow f to ensure the latter property as follows. Consider a pair $(u, v) \in \mathcal{M}$ such that $x'(u) \neq x'(v)$, and suppose without loss of generality that $x'(u) < x'(v)$. We create $x'(v)$ units of flow from u to z as follows: since $x'(v) \leq x(v)/c$, we can route $x'(v)$ units of flow from u to v using the flow paths of f and then route $x'(v)$ units of flow from v to z using the flow paths of g'. Therefore $f + 2g'$ contains a congestion three sub-flow that routes $x'(u) = x'(v)$ units of flow for each pair $(u, v) \in \mathcal{M}$. If we scale down this flow by a factor of 3, we obtain a feasible flow h with the desired properties. □

We preprocess the tree decomposition as in Section 3. Thus we may assume that G is connected and no bag is empty.

We prove Theorem 3.5 by induction on $\ell_1 + \ell_2 + |V(G)|$.

Base case. In the base case, we assume that $\ell_1 = \ell_2 = 0$. Since every parent adhesion of a non-root node is non-empty, that implies that the only bad node is the root t_0, that is, every flow path in f passes through $\beta(t_0)$, which is of size at most r.

We can obtain a flow g from $\mathcal{M}$ to $\beta(t_0)$ as follows: for each path P in the support of f, let P' be the smallest prefix of P that ends at a vertex of $\beta(t_0)$, and let $g(P') = f(P)$ (since each path P in the support of f intersects $\beta(t_0)$, there exists such a prefix P'). Using the flows f and g, we can construct a flow h in G from $V(\mathcal{M})$ to a single vertex $z \in \beta(t_0)$ with marginals π such that $\pi(u) = \pi(v)$ for each pair $(u, v) \in \mathcal{M}$ and $|h| \geq \frac{1}{3r}|f|$. (This follows from the argument used in the proof of Lemma 5.5.). Thus (G, π) is the desired decomposition.

In the inductive step, we consider two cases, depending on whether $0 \leq \ell_1 < \ell_2$ or $0 < \ell_1 = \ell_2$. The latter case can be handled in the same way as in Section 3. The former case requires a different argument.

Inductive step when $0 \leq \ell_1 < \ell_2$. Let $\{t_1, t_2, \ldots, t_p\}$ be the topmost bad nodes of $\mathcal{T}$ with parent adhesions of size ℓ_2, that is, it is a minimal set of such bad nodes such that for every bad node t with parent adhesion of size ℓ_2, there exists $1 \leq i \leq p$ with $t \preceq t_i$. For $1 \leq i \leq p$, let f_i^{inside} be the subflow of f consisting of all paths that are completely contained in $G[\alpha(t_i)]$. Furthermore, since $\ell_1 < \ell_2$, the node t_i is safe; let g_i be the corresponding flow, i.e., a flow that routes $\frac{1}{4r}x(v)$ from every $v \in \gamma(t_i)$ to $\sigma(t_i)$ in $G(t_i)$. By applying Lemma 5.5, for each i such that $1 \leq i \leq p$, there is a flow h_i that is nice for t_i and satisfies $|h_i| \geq \frac{1}{12r^2}|f_i^{\text{inside}}|$.

If $\sum_{i=1}^{p} |f_i^{\text{inside}}| > \frac{1}{r}|f|$, we construct the desired decomposition as follows. Let $z_i \in \sigma(t_i)$ be the endpoint of all the flow paths of h_i. Let $H_i = G[\alpha(t_i)] \cup \{z_i\}$. Note that h_i is completely contained in H_i and the graphs H_i are node-disjoint except for the vertices z_i. We group together the graphs with the same z_i vertex as follows. For each distinct vertex $z \in \{z_1, \ldots, z_p\}$, let $H_z = \bigcup \{H_i \colon 1 \leq i \leq p, z_i = z\}$ and $h_z = \sum \{h_i \colon 1 \leq i \leq p, z_i = z\}$. Note that h_z is a feasible flow in H_z whose flow paths all end at z. The desired decomposition has a component (H_z, π_z) for each distinct vertex z, where π_z are the marginals of h_z.

Therefore we may assume that $\sum_{i=1}^{p} |f_i^{\text{inside}}| \leq \frac{1}{r}|f|$. In this case, we drop the flows f_i^{inside}, that is, consider a flow $f' \coloneqq f - \sum_{i=1}^{p} f_i^{\text{inside}}$. Clearly, $|f'| \geq (1 - \frac{1}{r})|f|$. Furthermore, by the definition of f_i^{inside}, every node t_i is good with respect to f'. Since deleting a flow path cannot turn a good node into a bad one nor a safe node into an unsafe one, and all descendants of a good node are also good, we infer that every unsafe node with respect to f' has parent adhesion of size at most ℓ_1, while every bad node with respect to f' has parent adhesion of size *less* than ℓ_2. Consequently, by induction hypothesis we obtain a decomposition $(G_1, \pi_1), \ldots, (G_q, \pi_q)$ satisfying

$$\sum_{i=1}^{q} \pi_i(X_i) \geq \frac{1}{12r^3}\left(1 - \frac{1}{r}\right)^{\ell_1+\ell_2-1} |f'| \geq \frac{1}{12r^3}\left(1 - \frac{1}{r}\right)^{\ell_1+\ell_2} |f|.$$

Inductive step when $0 < \ell_1 = \ell_2$. This case follows from the argument given in Section 3. We define two node-disjoint instances $\mathcal{I}_1$ and $\mathcal{I}_2$ exactly as before. We apply the induction hypothesis to each of the instaces and obtain a decomposition for each. Since the instances are node-disjoint, the union of the two decompositions is the desired decomposition.

This concludes the proof of Theorem 5.3, which immediately implies Theorem 5.2. □

6 Hardness for MaxNDP in bounded treedepth graphs

In this section we prove Theorem 1.4.

We reduce from the $W[1]$-hard Multicolored Clique problem [15], where given a graph G, an integer k, and a partition $V = V^1 \uplus V^2 \uplus \ldots \uplus V^k$, we are to check if there exists k-clique in G with exactly one vertex in every set V^i. By adding dummy vertices, we can assume that $|V^i| = n$ for every $1 \leq i \leq k$, and that $n, k \geq 2$.

Construction. Given an instance $(G, k, (V^i)_{i=1}^k)$ of Multicolored Clique, we aim at constructing an equivalent instance $(H, \mathcal{M}, \ell)$ of MaxNDP.

We start with a construction, for every set V^i, a gadget W^i as follows. First, for every $v \in V^i$ we construct an $(k-1)$-vertex path X^i_v on vertices $x^i_{v,1}, x^i_{v,2}, \ldots, x^i_{v,i-1}, x^i_{v,i+1}, \ldots, x^i_{v,k}$. Second, we select an arbitrary vertex $u^i \in V_i$. Third, for every $v \in V^i \setminus \{u^i\}$, we add a vertex s^i_v adjacent to the first vertex of X^i_v (i.e., $x^i_{v,1}$ and $x^i_{u^i,1}$ if $i > 1$ or $x^i_{v,2}$ and $x^i_{u^1,2}$ if $i = 1$), a vertex t^i_v adjacent to the last vertex of X^i_v (i.e., $x^i_{v,k}$ and $x^i_{u^i,k}$ if $i < k$ or $x^i_{v,k-1}$ and $x^i_{u^i,k-1}$ if $i = k$), and make (s^i_v, t^i_v) a terminal pair. This concludes the description of the gadget W^i. By $\mathcal{M}_{st}$ we denote the set of terminal pairs constructed in this step.

To encode adjacencies in G, we proceed as follows. For every pair $1 \leq i < j \leq k$, we add a vertex $p_{i,j}$, adjacent to all vertices $x^i_{v,j}$ for $v \in V_i$ and all vertices $x^j_{u,i}$ for $u \in V_j$. For every edge $vu \in E(G)$ with $v \in V_i$ and $u \in V_j$, we add a terminal pair $(x^i_{v,j}, x^j_{u,i})$. Let $\mathcal{M}_x$ be the set of terminal pairs constructed in this step; we have $\mathcal{M} = \mathcal{M}_{st} \cup \mathcal{M}_x$.

Finally, we set the required number of paths $\ell := k(n-1) + \binom{k}{2}$. This concludes the description of the instance $(H, \mathcal{M}, \ell)$; see Figure 3 for an illustration.

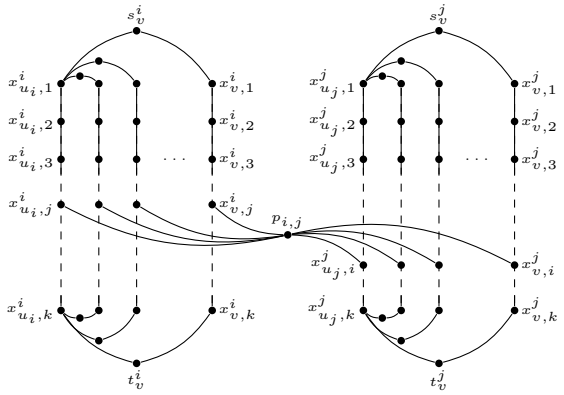

Figure 3: Part of the construction of the graph H: gadgets W^i and W^j, together with a connection via $p_{i,j}$.

From a clique to disjoint paths. Assume that the input Multicolored Clique instance is a "yes"-instance, and let $\{v^i : 1 \leq i \leq k\}$ be a clique in G with $v^i \in V^i$ for every $1 \leq i \leq k$. We construct a family of ℓ vertex-disjoint paths as follows. First, for every $1 \leq i \leq k$ and every $v \in V^i \setminus \{u^i\}$, we route a path from s^i_v to t^i_v through the path X^i_v if $v \neq v^i$, and through the path $X^i_{u^i}$ if $v = v^i$. Note that in this step we have created $k(n-1)$ vertex-disjoint paths connecting terminal pairs, and in every gadget W^i the only unused vertices are vertices on the path $X^i_{v^i}$. To construct the remaining $\binom{k}{2}$ paths, for every pair $1 \leq i < j \leq k$ we take the 3-vertex path from $x^i_{v^i,j}$ to $x^j_{v^j,i}$ through $p_{i,j}$; note that the assumption that $v^i v^j \in E(G)$ ensures that $(x^i_{v^i,j}, x^j_{v^j,i})$ is indeed a terminal pair in $\mathcal{M}$.

From disjoint paths to a clique. In the other direction, let $\mathcal{P}$ be a family of ℓ vertex-disjoint paths connecting terminal pairs in H. Let $\mathcal{P}_{st} \subseteq \mathcal{P}$ be the set of paths connecting terminal pairs from $\mathcal{M}_{st}$, and similarly define $\mathcal{P}_x$. First, observe that the set $P := \{p_{i,j} : 1 \leq i < j \leq k\}$ separates every terminal pair from $\mathcal{M}_x$. Hence, every path from $\mathcal{P}_x$ contains at least one vertex from P. Since $|P| = \binom{k}{2}$, we have $|\mathcal{M}_x| \leq \binom{k}{2}$, and, consequently, $|\mathcal{P}_{st}| \geq \ell - \binom{k}{2} = k(n-1) = |\mathcal{M}_{st}|$. We infer that $\mathcal{P}_{st}$ routes all terminal pairs in $\mathcal{M}_{st}$ without using any vertex of P, while $\mathcal{P}_x$ routes $\binom{k}{2}$ pairs from $\mathcal{M}_x$, and every path from $\mathcal{P}_x$ contains exactly one vertex from P.

Since the paths in $\mathcal{P}_{st}$ cannot use any vertex in P, every such path needs to be contained inside one gadget W^i. Furthermore, observe that a shortest path between terminals $s^i_{v,a}$ and $t^i_{v,a}$ inside W^i is either $X^i_{u^i}$ or X^i_v, prolonged with the terminals at endpoints, and thus contains $k+1$ vertices. Furthermore, a shortest path between two terminals in $\mathcal{M}_x$ contains three vertices. We infer that the total number of vertices on

paths in $\mathcal{P}$ is at least

$$|\mathcal{P}_{st}| \cdot (k+1) + |\mathcal{P}_x| \cdot 3 = k(n-1)(k+1) + 3\binom{k}{2} = k\left(n(k-1) + 2(n-1)\right) + \binom{k}{2} = |V(H)|.$$

We infer that every path in $\mathcal{P}_{st}$ consists of $k+1$ vertices, and every path in $\mathcal{P}_x$ consists of three vertices. In particular, for every $1 \leq i \leq k$ and $v \in V^i \setminus \{u^i\}$, the path in $\mathcal{P}_{st}$ that connects s^i_v and t^i_v goes either through X^i_v or $X^i_{u^i}$. Consequently, for every $1 \leq i \leq k$ there exists a vertex $v^i \in V^i$ such that the vertices of W^i that do not lie on any path from $\mathcal{P}_{st}$ are exactly the vertices on the path $X^i_{v^i}$.

We claim that $\{v^i : 1 \leq i \leq k\}$ is a clique in G. To this end, consider a pair $1 \leq i < j \leq k$. Since $|\mathcal{P}_x| = \binom{k}{2}$, there exists a path in $\mathcal{P}_x$ that goes through $p_{i,j}$. Moreover, this path has exactly three vertices. Since the only neighbours of $p_{i,j}$ that are not used by paths from $\mathcal{P}_{st}$ are $x^i_{v^i,j}$ and $x^j_{v^j,i}$, we infer that $(x^i_{v^i,j}, x^j_{v^j,i}) \in \mathcal{M}$ and, consequently, $v^i v^j \in E(G)$. This concludes the proof of the correctness of the construction.

Treedepth bound. We are left with a proof that H has bounded treedepth. To this end, we use the alternative definition of treedepth [21]: treedepth of an empty graph is 0, while for any graph G on at least one vertex we have that

$$\mathrm{td}(G) = \begin{cases} 1 + \min\{\mathrm{td}(G-v) : v \in V(G)\} & \text{if } G \text{ is connected} \\ \max\{\mathrm{td}(C) : C \text{ connected component of } G\} & \text{otherwise.} \end{cases}$$

First, observe that $H - P$ contains k connected components, being the gadgets W^i. Second, observe that the deletion of the endpoints of the path $X^i_{u^i}$ from the gadget W^i breaks W^i into connected components being paths on at most $k+1$ vertices. Consequently,

$$\mathrm{td}(H) \leq |P| + 2 + k + 1 = \mathcal{O}(k^2).$$

This finishes the proof of Theorem 1.4. □

References

[1] A. Andoni, A. Gupta, and R. Krauthgamer. Towards $(1+\varepsilon)$-approximate flow sparsifiers. In *Proc. SODA 2014*, pages 279–293, 2014.

[2] M. Andrews, J. Chuzhoy, V. Guruswami, S. Khanna, K. Talwar, and L. Zhang. Inapproximability of edge-disjoint paths and low congestion routing on undirected graphs. *Combinatorica*, 30(5):485–520, 2010.

[3] P. Chalermsook, J. Chuzhoy, A. Ene, and S. Li. Approximation algorithms and hardness of integral concurrent flow. In *Proc. STOC 2012*, pages 689–708, 2012.

[4] C. Chekuri and J. Chuzhoy. Large-treewidth graph decompositions and applications. In *Proc. STOC 2013*, pages 291–300, 2013.

[5] C. Chekuri and A. Ene. Poly-logarithmic approximation for maximum node disjoint paths with constant congestion. In *Proc. SODA 2013*, pages 326–341, 2013.

[6] C. Chekuri, S. Khanna, and F. Shepherd. An $O(\sqrt{n})$ approximation and integrality gap for disjoint paths and unsplittable flow. *Theory Comput.*, 2(7):137–146, 2006.

[7] C. Chekuri, S. Khanna, and F. B. Shepherd. Multicommodity flow, well-linked terminals, and routing problems. In *Proc. STOC 2005*, pages 183–192, 2005.

[8] C. Chekuri, S. Khanna, and F. B. Shepherd. Edge-disjoint paths in planar graphs with constant congestion. *SIAM J. Comput.*, 39(1):281–301, 2009.

[9] C. Chekuri, S. Khanna, and F. B. Shepherd. A note on multiflows and treewidth. *Algorithmica*, 54(3):400–412, 2009.

[10] C. Chekuri, G. Naves, and F. B. Shepherd. Maximum edge-disjoint paths in k-sums of graphs. In *Proc. ICALP 2013*, volume 7965 of *LNCS*, pages 328–339, 2013.

[11] E. Chlamtac, R. Krauthgamer, and P. Raghavendra. Approximating sparsest cut in graphs of bounded treewidth. In *Proc. APPROX-RANDOM 2010*, volume 6302 of *LNCS*, pages 124–137, 2010.

[12] J. Chuzhoy. On vertex sparsifiers with Steiner nodes. In *Proc. STOC 2012*, pages 673–688, 2012.

[13] J. Chuzhoy and S. Li. A polylogarithmic approximation algorithm for edge-disjoint paths with congestion 2. In *Proc. FOCS 2012*, pages 233–242, 2012.

[14] R. Diestel. *Graph theory*, volume 173. Springer, third edition, 2005.

[15] M. R. Fellows, D. Hermelin, F. Rosamond, and S. Vialette. Parameterized complexity of multiple-interval graph problems. *Theoret. Comput. Sci.*, 410(1):53–61, 2009.

[16] N. Garg, V. V. Vazirani, and M. Yannakakis. Primal-dual approximation algorithms for integral flow and multicut in trees. *Algorithmica*, 18(1):3–20, 1997.

[17] A. Gupta, K. Talwar, and D. Witmer. Sparsest cut on bounded treewidth graphs: algorithms and hardness results. In *Proc. STOC 2013*, pages 281–290, 2013.

[18] R. M. Karp. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, pages 85–103, 1972.

[19] S. Kolliopoulos and C. Stein. Approximating disjoint-path problems using packing integer programs. *Math. Prog.*, 99(1):63–87, 2004.

[20] T. Leighton and S. Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *J. ACM*, 46(6):787–832, 1999.

[21] J. Nešetřil and P. O. de Mendez. Tree-depth, subgraph coloring and homomorphism bounds. *Eur. J. Comb.*, 27(6):1022–1041, 2006.

[22] J. Nešetřil and P. O. de Mendez. *Sparsity - Graphs, Structures, and Algorithms*, volume 28 of *Algorithms and combinatorics*. Springer, 2012.

[23] N. Robertson and P. D. Seymour. Graph minors. V. Excluding a planar graph. *J. Combinatorial Theory, Ser. B*, 41(1):92–114, 1986.

[24] N. Robertson and P. D. Seymour. Graph minors. XIII. The disjoint paths problem. *J. Combinatorial Theory, Ser. B*, 63(1):65–110, 1995.

[25] N. Robertson and P. D. Seymour. Graph minors. XVI. Excluding a non-planar graph. *J. Combinatorial Theory, Ser. B*, 89(1):43–76, 2003.

[26] L. Séguin-Charbonneau and F. B. Shepherd. Maximum edge-disjoint paths in planar graphs with congestion 2. In *Proc. FOCS 2011*, pages 200–209, 2011.